

CHAPTER 4

DESIGNING FOR RELIABILITY

4-1. Establish and allocate requirements

For a new product or system, developing requirements is the first step, whether the requirement is reliability or any other performance characteristic. Requirements must be realistic. They should be derived from the customer's or user's needs (the mission), economic considerations (life cycle cost), and other factors. For guidance in addressing the reliability and availability of C4ISR facilities during design and in operation, see TM 5-698-1.

a. *Deriving requirements.* Many ways of deriving reliability requirements are used. Some are based on achieving incremental improvements in each successive model of a product. Others are derived from sophisticated simulations that model the way in which the system will be used. Still others, benchmarking for example, are based on staying competitive with other suppliers. It is important to note that customers often state reliability requirements in a way that is not directly usable by designers. Also, designers do not always have direct control over all of the factors that influence the reliability that will be achieved in use.

(1) Customers and system users often think not of reliability, but of availability – how often the system will be available for use – or a maximum number of warranty returns. It is difficult for designers to work directly with these types of requirements. Consequently, a "translation" must be made to convert these higher-level requirements to design measures, such as probability of failure or MTBF. For example, if availability is the customer's requirement, many combinations of reliabilities and repair times will result in the required availability.

(2) The reliability achieved for a system in use is affected not only by the design and manufacturing processes, but also by the skill and experience of the operators and maintainers, and by changes in the way the system is operated. Designers may not be able to control all of these factors. For example, designers can consciously attempt to minimize the possibility of failures being induced during maintenance but cannot prevent all such failures from occurring. However, the design requirement can be "adjusted" so that even with some reasonable number of maintenance-induced failures, the reliability in actual use will meet the customer's needs. This adjustment means that the design requirement must be higher than one would first imagine.

b. *Allocating requirements.* Customers and users usually state the reliability requirement (or a high-level requirement having reliability as a key element) at the product or system level. For example, the reliability for an electrical power generation system might be 99.9% for a given power level into a given load for a stated period of time. But what should be the reliability requirement for a transformer used in the system?

(1) To better understand the reliability allocation process, consider how weight is treated. If a maximum weight is specified for a system, each element of the system must be assigned a weight "budget" that the designers must meet. If, for example, a system consists of 5 elements A through E and the system weight must be no more than 2,000 lbs., we might assign budget as follows: A - 200 lbs., B - 500 lbs., C - 350 lbs., D - 400 lbs., and E - 550 lbs. The sum of the element weights must, of course, add up to no more than the maximum system weight. The assignment of the budgets would be made on past experience or some other logical basis.

(2) The allocation of a system reliability requirement is similar to the assignment of weight budgets. The idea is to assign reliability requirements to lower levels of indenture within the system such that if the lower-level requirements are met, the system requirement will be met. For example, if the system reliability for a 10-hour mission is specified as 95%, and the system is made up of three major subsystems A, B, and C, then $R_A \times R_B \times R_C$ must be equal to 0.95.

(3) Several methods are used to make reliability allocations. These include the Equal Allocation Method, the ARINC Method, and the Feasibility of Objectives Method. These and other methods are described in several of the references listed in appendix A.

4-2. Develop system reliability model

Early in the development of a new system, a reliability model must be developed. The most commonly used model is the reliability block diagram (RBD) discussed in chapter 3. The process for modeling a system for reliability purposes consists of three steps.

- a. *Select system.* Define the specific system to be modeled. This definition includes the exact configuration and, if appropriate, the block or version.
- b. *Construct functional block diagram.* The functional relationships among the parts, assemblies, and subsystems must be understood because reliability deals with functional failures. In fact, failure is usually defined as the loss of a function. The functional block diagram shows inputs and outputs but does not necessarily depict how the system elements are physically connected or positioned. Figure 4-1 shows an example of a functional block diagram.

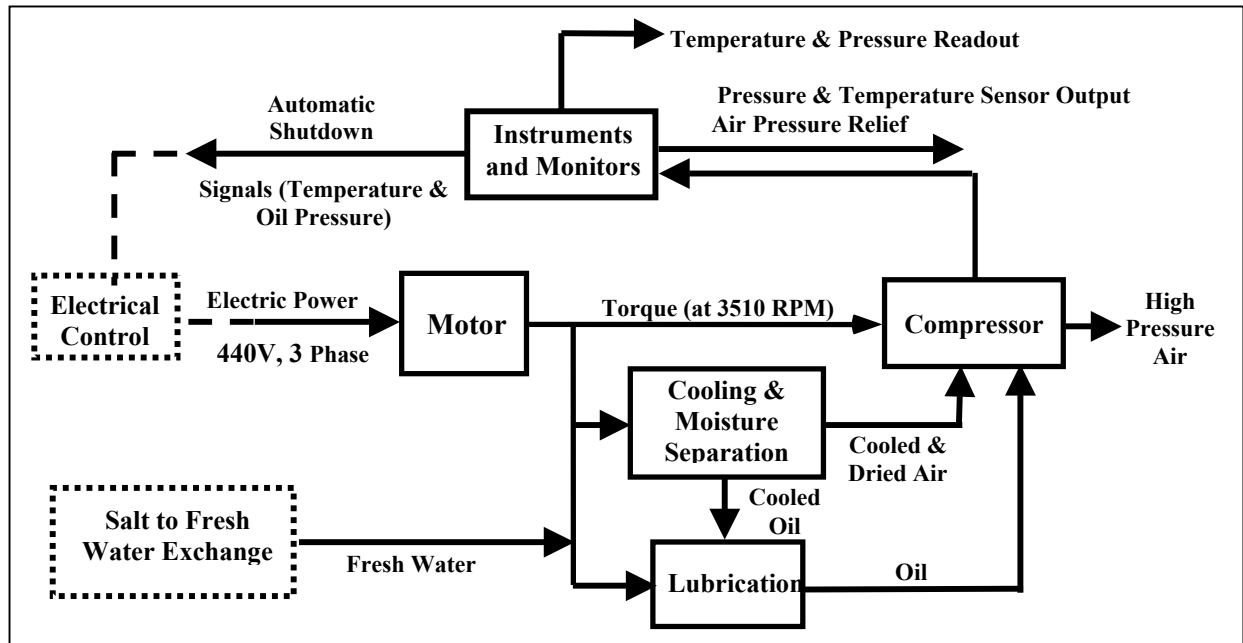


Figure 4-1. Example of a functional block diagram.

- c. *Construct reliability block diagrams as necessary.* It is often impractical to develop one RBD for the entire system that has all subsystems, assemblies, and parts. A single RBD for an entire C4ISR facility would be huge and unmanageable. More commonly, RBDs are developed for lower-level portions of the system, such as the subsystem, assembly, and even part level. The reliability of each of these portions can then be assessed and used in a system assessment. Figure 4-2 illustrates this process.

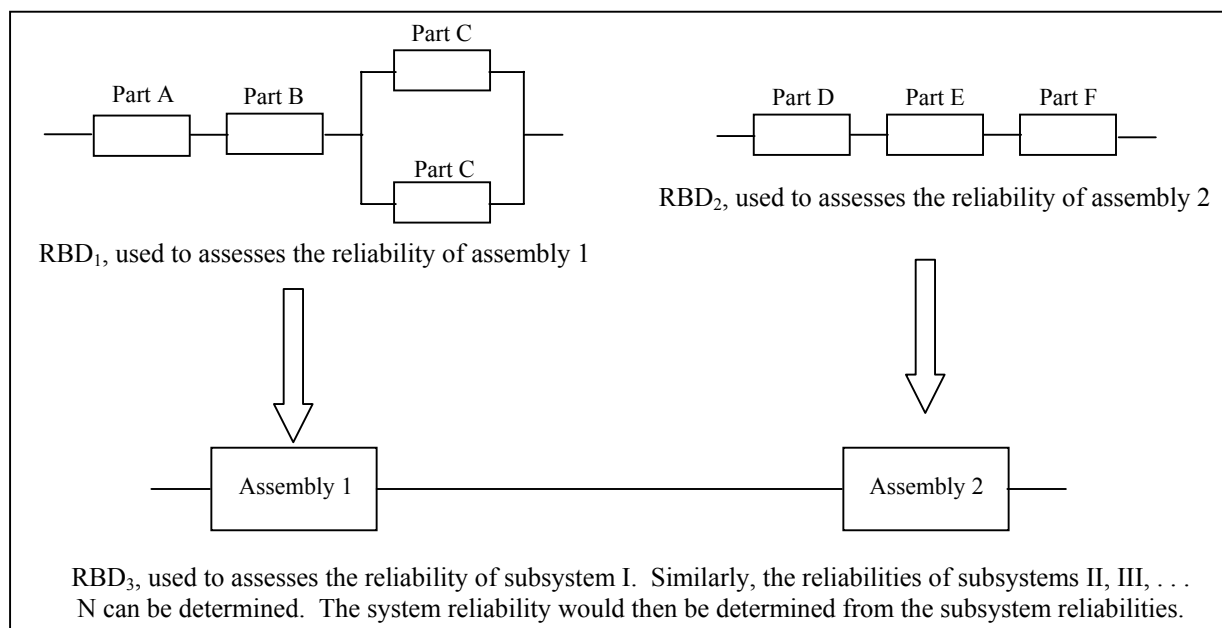


Figure 4-2. An example of how lower-level RBDs are used to assess the reliabilities of assemblies. The resulting assembly reliabilities are used in an RBD of a subsystem made up of the assemblies. This process can be repeated until the system reliability can be assessed.

4-3. Conduct analyses

A variety of analyses can be used in designing for reliability. Table 4-1 lists the titles and purposes of some of these analyses.

a. *Related analyses.* Many analyses are conducted for reasons not specifically stated as reliability, such as safety and structural integrity. However, many of these analyses directly or indirectly support the effort of designing for reliability. Designers should always have the objective of using the results of analyses for as many purposes as practical. An integrated systems approach facilitates extracting as much benefit from all analyses (as well as tests).

Table 4-1. Typical reliability-related analyses and their purposes

Analysis	Purpose
Dormancy Analysis	Used to calculate failure rates of devices while dormant (e.g., storage).
Durability Assessment	Used to confirm a design life for a product. It is more effectively applied earlier in development to ensure that design life is adequate.
Failure Modes, Effects, and Criticality Analysis	Used ideally as a design and assessment tool to understand and alleviate failure consequences, it can also be an independently applied tool to check that certain failure consequences are avoided. A qualitative measurement.
Fault Tree Analysis (FTA)	Used ideally as a design and assessment tool to understand and alleviate failure consequences, it can also be an independently applied tool to check that certain failure consequences are avoided. A qualitative measurement.
Finite Element Analysis (FEA)	FEA is a computer simulation technique used for predicting material response or behavior of modeled device, determining material stresses and temperature, and determining thermal and dynamic loading.
Sneak Circuit Analysis (SCA)	Used ideally as a design and assessment tool to discover unintended paths and functions, it can also be an independently applied tool to check that certain failure consequences are avoided. A qualitative measurement.
Thermal Analysis (TA)	Used to calculate junction temperatures, thermal gradients, and operating temperatures.
Worst Case Circuit Analysis (WCCA)	A tool used to effectively assess design tolerance to parameter variation, it can also be used as an independent check of the susceptibility to variation.

b. *The Role of the designer.* In some cases, designers will and should be directly involved in performing a given analysis. Other individuals may perform specific and highly specialized analyses. In any case, it is important that the designers understand the purpose and benefit of each analysis, and "buy in" to the need for conducting the analysis.

4-4. Design for reliability

Achieving the required level of reliability begins with design. Some key issues that must be addressed during design are control of parts and materials, use of redundancy, robust design, design from the environment, designing for simplicity, and configuration control.

a. *Control selection of parts and materials.* Part of the design for reliability process is the selection of parts and materials. In selecting parts and materials, the designers must consider functionality, performance, reliability, quality, cost, producibility, long-term availability, and other factors.

(1) When possible, standard parts and materials having well-established characteristics should be preferred to non-standard or newly developed parts and materials. For some products or use environments, the anticipated stresses are so low that any commercially available part may be acceptable. In such cases, parts control may consist entirely of configuration management (knowing what parts are used) and ensuring that they are obtained from a reputable source. In other cases, the stresses that will be encountered by the product may eliminate many types of parts or mandate certain application criteria (e.g., derating). In addition, some types of parts may be obsolete before the product is delivered. In these cases, parts control should be more extensive and rigorous.

(2) After selecting the appropriate part it should be applied in a conservative manner (a process called derating). Using a part at its maximum capability increases the failure rate and does not allow for transients or overloads. Just how conservatively a part may be used depends on factors such as cost, mission criticality, and environment, which cannot be generalized.

b. *Use redundancy appropriately.* You will recall that components or subsystems connected in parallel must all fail in order to have system failure. This addition of components or subsystems in parallel is termed redundant configurations. Simply stated, redundancy provides alternate paths or modes of operation that allow for proper system operation. Redundancy has some drawbacks, however, and cannot be blindly used. Adding parallel items increases weight and cost. It increases complexity. Finally, redundancy does nothing to increase the reliability of individual items, only the system-level mission reliability. It actually decreases basic reliability. Thus, more failures (albeit not mission failures) will occur requiring repair or replacement, driving up support costs.

c. *Use robust design.* A robust system design is one that is tolerant of failures and occasional spikes in stresses. One way to achieve a robust design is to use Design of Experiments to determine which parameters are critical and then to optimize those parameters. Another method involves the use of Highly Accelerated Life Testing (HALT). HALT requires successively higher stresses to be applied during test and making design changes to eliminate the failures observed at each level of stress. The magnitude of the stresses is not intended to represent actual use but to force failures. Using HALT results in "over-designed" systems and products, but over-design may be warranted in critical applications.

d. *Design for the environment.* Without an understanding of the environment to which a system will be exposed during its useful life, designers cannot adequately design for or predict reliability. The process of understanding a system's environment is referred to as environmental characterization. The environment includes not only the operating environment but also all other environments applicable to the system. Often, the operating environment does not impose the greatest stresses. Table 4-2 lists some of the environments that must be considered in designing for reliability.

Table 4-2. Environments to consider in designing for reliability

Environment	Comments	Environmental Stresses and Factors*
Operating	Includes all potential ways and climates in which the system will be used.	Temperature Humidity Mechanical/acoustical vibration Mechanical/acoustical shock Moisture Sand Dirt Electromagnetic interference Radiation Mechanical loads Corrosion Chemical reaction
Support	The environment in which a system is repaired and serviced must be considered.	
Installation	For some systems, the process of installation imposes stresses that are higher than those of operation.	
Storage	For systems and products stored for long periods of time, the storage environment can be the dominant cause of failure.	
Transportation	The shipping and handling of systems and products can impose stresses, such as shock and vibration, that are different from or higher than those of operation.	

*Typical environmental stresses and factors that can occur in any of the listed environments.

e. *Design for simplicity.* The basic tenet of reliable design is to keep it simple. The more complicated a design, the more opportunity for failure. This principle is sometimes derided as elementary and intuitive; nevertheless, it is often needlessly violated and is included here as a reminder of its importance.

f. *Institute configuration control.* As changes are made to improve reliability, or for any other reason, and the design matures, it should be complemented by a progressively mature control of hardware design. It is important to know which current configuration served as the basis for a given reliability prediction or analysis.

(1) Initially, the hardware design is conceptual in nature and may be described by equations or design parameters, for example. At this stage, subsystem designers should have little controls placed upon the details. They should be engaged in trade studies, sensitivity analyses, and design variations leading to the next phase of hardware control.

(2) The next level of configuration control is "baselining the system." The baseline permits concentration on a specific design and allows detail design to begin. After a system is baselined, the designer can only change the concept when there is due cause and only after notifying other program elements to assure that each subsystem designer is aware of the design of interfacing subsystems.

(3) At critical design review (drawing release), the detail design is (ideally) complete and formal configuration control process should be instituted. The process should be rigid and designed to ensure that design modifications are undertaken only for understood cause and the full cost and impact is analyzed prior to initiating the change.

4-5. Conduct development testing

Reliability prediction and design requires some knowledge of the failure rates of parts, and how the parts are used. Additionally, the reliability engineer will need to use analytical tools such as FMEAs and stress analysis. In performing analyses and making predictions, the engineer tries to account for all factors affecting reliability. However, as is true of all analysis, the reliability analysis is far from perfect, particularly early in the development of a new product. For instance, initial tests of the product (the product may be a prototype, development model, or production article) may reveal unforeseen failure modes. Then again, it might be determined that initial failure rates and application factors did not sufficiently account for interaction of parts and subsystems (the fact that the whole is not always the simple sum of its parts is attributed to a phenomenon called synergism). Consequently, the MTBF (hardware reliability) or mission reliability may be lower than originally estimated. Since the original design was intended to satisfy a requirement, some action is needed to bring the reliability of the product "up to spec." The process by which the reliability of a product is increased to meet the design level is reliability growth.

a. *Duane's model.* Duane developed learning curves based on cumulative failures and cumulative operating hours for five different products: two hydromechanical devices, two aircraft electrical generators, and a turbojet

engine. The products represented a broad range of aircraft type equipment and were identified only by general description. After plotting the data on log-log paper, Duane found that the curves were very nearly linear and that failure rates at any point in time for these relatively complex aircraft accessories were approximately inversely proportional to the square root of the cumulative operating time. Independent and related efforts such as the Godovin Report, work by J.D. Selley, S.G. Miller, and E.O. Codier of General Electric, and others have confirmed the soundness of Duane's hypothesis. In total, this work has given the engineer and the manager an aid in planning, monitoring, and controlling the growth of reliability early in an acquisition program.

b. *Other Models for Growth.* Duane's work has been expanded and extended by engineers and statisticians and a variety of reliability growth models are now available. One, the AMSAA-Crow model is a statistical model based on the Non-Homogeneous Poisson Process (NHPP). The NHPP applies when a trend exists (e.g., reliability is improving or degrading). Since the AMSAA-Crow is a statistical model, it is somewhat more complicated to use than the Duane model.

(1) First, you must determine if a trend exists in the data using a statistic called the Laplace statistic (this statistic will be addressed in more detail in chapter 6). If a trend does not exist at some level of confidence, determined by the user, the model cannot be used.

(2) If the model applies, then you calculate parameters based on sample size and type of test (test ended after a given number of failures or after a given length of time).

(3) You can now determine system failure rate at time of interest

(4) An advantage of the AMSAA-Crow model is that, since it is a statistical model, you can calculate upper and lower bounds on calculated failure rate (or the MTBF).

c. *Achieving reliability growth.* Corrective measures taken to ensure that the equipment reliability "grows" properly include redesign, change in materials or processes, or increased tolerances on critical parameters. All of these efforts represent the expenditure of money.

d. *The nature of growth.* Reliability growth is the decrease in the hazard function during the early portion of development and production. It is the result of design changes and improvements that correct deficiencies of the original design. Its goal is to attain a design which, when in full operational use, has the minimum required level of reliability. When reliability growth is completed, the hazard function (failure rate if the exponential distribution applies) stabilizes at a relatively fixed value. The key attributes of reliability growth follow.

(1) Reliability growth occurs early in the life cycle of a product.

(2) Reliability growth is the result of corrective action. Reliability growth is intended to achieve the required reliability level. Testing provides verification of the predictions made as a result of analytical methods and of the design approach used. When testing reveals that the analyses or design approaches were inadequate or deficient, corrective actions must be taken to the extent necessary to meet the reliability requirements. Assuming the corrective actions are effective, growth occurs.

(3) The hazard function stabilizes when growth ceases. For systems, which tend to exhibit times between failure that are exponentially distributed, this behavior means that once growth ceases, we will observe a constant failure rate (or a constant mean time between failure). The value will, of course, actually fluctuate due to variances in operations and other factors but will be relatively stable. When the system starts to near the end of its useful life, the failure rate will start to increase. Trending is intended to provide an early indication when system reliability is degrading (due to age or for other reasons).

e. *Accelerated testing.* Earlier, HALT was introduced as a technique for achieving a robust design. HALT is one form of accelerated testing. Another, Accelerated Life Testing (ALT), is a technique for achieving reliability growth by accelerating the rate at which failures occur and are addressed by design improvements. The primary difference between HALT and ALT is that the accelerated stresses used in the latter are chosen such that failures not expected in actual use (storage, installation, etc.) are hopefully not introduced during the ALT. This constraint allows the

results of ALT to be used to assess the reliability of the item being tested. HALT does not provide an estimate of the true reliability, only some assurance that the reliability is higher than some minimum.

4-6. Iterate the design

As discussed briefly in paragraph 4-5, as changes are made to improve reliability, or for any other reason, the design is changed and gradually matures. This iteration process is an inherent part of design and development, especially when the system is new or significantly different from predecessors. Changes to the design are made on the basis of continuing analyses. These analyses are initially performed on conceptual designs and eventually on test results. When these changes are made to reduce the relative frequency with which a failure occurs, or to reduce or minimize the effects of a failure, the change is related to reliability growth.

4-7. Conduct demonstration tests

At or near the end of development, a key question that must be answered is "Has the reliability requirement been met?" Either the customer will require that some type of test be made to measure the level of reliability achieved or the company itself may require such testing as a matter of policy. Such tests are called demonstration tests because they are intended to demonstrate the level of reliability that has actually been achieved. Ideally, such testing would be conducted on products right off the production line. Practical considerations make this nearly impossible. Indeed, the decision whether or not to proceed with full-scale production often requires the testing to have been completed. Consequently, testing is done using early production models or prototypes that are as close as practical to the full-rate production model.

a. *Standard statistical tests.* For many years, the statistical tests described in MIL-HDBK-781 were used to demonstrate the achieved level of reliability. MIL-HDBK-781 provides for two types of tests: sequential tests (called probability ratio sequential tests) and fixed-length tests. Both types are based on the premise that a product (system) exhibits a constant failure rate (i.e., the underlying pdf is the exponential) and is neither getting more reliability or less reliable. The problem with such tests is that for products having high MTBFs, the test time can be very long. For example, a fixed-length test to verify that a product has an MTBF of 1,000 hours can take as many as 45,000 hours of cumulative test hours. If a sample of only 3 is available, that means the test will take 15,000 calendar hours. Such testing is obviously expensive.

b. *Accelerated life testing.* Accelerated life testing was introduced in Paragraph 4-5e as a technique for accelerating reliability growth. It can also be used to avoid the problem of demonstrating very high reliabilities with MIL-HDBK-781 tests and is finding an increasingly larger following for this purpose. Accelerated testing is intended to accelerate the occurrence of failures that would eventually occur under normal conditions, measure the reliability at these conditions, and then correlate the results back to normal operating stresses. In accelerating the stresses, it is important not to induce failures that would not otherwise occur. Otherwise, correlation is lost. Accelerated testing of complex systems has many uncertainties and not all failure modes can be accelerated.

4-8. Design using a team approach

The engineer and manager are both continually making trade-offs between complexity and flexibility, design costs and logistics support costs, redundancy and simplicity, etc. The ultimate goal of each member of the management and design team should be to obtain the essential operational performance characteristics at the lowest life cycle cost. To this end, the manager, engineer, logistic planner, and entire program team must maintain a daily dialogue, each contributing his talents to the benefit of all.

a. *Production and logistics affect reliability.* The designer works with ideas. Once these ideas were captured in hard-copy drawings and specifications; today they are captured in digital format. These ideas must be converted from the abstract to the concrete; i.e., from drawings and specifications to hardware. The process by which this conversion takes place is production. The manufacturing processes and the control of those processes affect the reliability of the finished system. Logistics also affects reliability.

(1) The manufacturing people can determine if the necessary processes, machines, skills, and skill levels are already in being. If not, they will have time to plan for these items (e.g., procure or develop new machines and

processes, hire new people, develop training, and so forth). The manufacturing people can help the designers by pointing out design approaches that are too complicated or impractical for manufacture. They can describe current capabilities to help designers select appropriate design approaches.

(2) The reliability observed by the customer is also affected by how well maintenance is performed and by the number of induced failures caused by inexperienced, careless, or inadequately trained personnel. The availability, even if the reliability being achieved in use is adequate, will be less than desired if the necessary trained people, spares, test equipment, or other logistics resources are unavailable when needed. Although availability will suffer, reliability is often incorrectly singled out as the problem.

b. *Everyone can benefit from the team approach.* Other people and organization who can contribute to the design for reliability and who can benefit from reliability analyses include the safety engineers, logistics planners, mission planners, packaging and handling specialists, after-sale service organizations, and so forth.

c. *Integrated Product and Process Development.* Within the Department of Defense, the Integrated Product and Process Development (IPPD) approach has been mandated for all DoD acquisition programs. IPPD is described in "DoD Guide to Integrated Product and Process Development" (version 1.0 was released February 5, 1996). Integrated Product Teams (IPTs) that organize for and accomplish the tasks required in acquiring goods and services are the foundation of the IPPD process. IPTs are made up of everyone having a stake in the outcome or product of the team, including the customer and suppliers. Collectively, team members should represent the needed know-how and be given the authority to control the resources necessary for getting the job done.

d. *The systems engineering approach.* Systems engineering is an interdisciplinary approach that focuses on defining customer needs and required functionality early in the development cycle, documenting requirements, then proceeding with design synthesis and system validation while considering the complete problem. The complete problem involves operations, performance, test, manufacturing, cost and schedule, training and support, and disposal. Systems engineering integrates all the disciplines and specialty groups into a team effort and promotes a structured development process that proceeds from concept to production to operation. Systems engineering is focused on the goal of providing a quality product that meets the user needs.